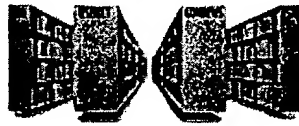


Internet RFC/STD/FYI/BCP Archives



RFC1338

[[Index](#) | [Search](#) | [What's New](#) | [Comments](#) | [Help](#)]

Network Working Group
Request for Comments: 1338

V. Fuller
BARRNet
T. Li
cisco
J. Yu
MERIT
K. Varadhan
OARnet
June 1992

Supernetting: an Address Assignment and Aggregation Strategy

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Abstract

This memo discusses strategies for address assignment of the existing IP address space with a view to conserve the address space and stem the explosive growth of routing tables in default-route-free routers run by transit routing domain providers.

Table of Contents

Acknowledgements	2
1. Problem, goal, and motivation	2
2. Scheme plan	3
2.1. Aggregation and its limitations	3
2.2. Distributed network number allocation	5
3. Cost-benefit analysis	6
3.1. Present allocation figures	7
3.2. Historic growth rates	8
3.3. Detailed analysis	8
3.3.1. Benefits of new addressing plan	9
3.3.2. Growth rate projections	9
4. Changes to Inter-Domain routing protocols	11
4.1. General semantic changes	11
4.2. Rules for route advertisement	11
4.3. How the rules work	13
4.4. Responsibility for and configuration of aggregation	14
5. Example of new allocation and routing	15
5.1. Address allocation	15
5.2. Routing advertisements	17
6. Transitioning to a long term solution	18

7. Conclusions	18
8. Recommendations	18
9. Bibliography	19
10. Security Considerations	19
11. Authors' Addresses	19

Acknowledgements

The authors wish to express their appreciation to the members of the ROAD group with whom many of the ideas contained in this document were inspired and developed.

1. Problem, Goal, and Motivation

As the Internet has evolved and grown over in recent years, it has become painfully evident that it is soon to face several serious scaling problems. These include:

1. Exhaustion of the class-B network address space. One fundamental cause of this problem is the lack of a network class of a size which is appropriate for mid-sized organization; class-C, with a maximum of 254 host addresses, is too small while class-B, which allows up to 65534 addresses, is too large to be widely allocated.
2. Growth of routing tables in Internet routers beyond the ability of current software (and people) to effectively manage.
3. Eventual exhaustion of the 32-bit IP address space.

It has become clear that the first two of these problems are likely to become critical within the next one to three years. This memo attempts to deal with these problems by proposing a mechanism to slow the growth of the routing table and the need for allocating new IP network numbers. It does not attempt to solve the third problem, which is of a more long-term nature, but instead endeavors to ease enough of the short to mid-term difficulties to allow the Internet to continue to function efficiently while progress is made on a longer-term solution.

The proposed solution is to hierarchically allocate future IP address assignment, by delegating control of segments of the IP address space to the various network service providers.

It is proposed that this scheme of allocating IP addresses be undertaken as soon as possible. It is also believed that the scheme will suffice as a short term strategy, to fill the gap between now

and the time when a viable long term plan can be put into place and deployed effectively. It is believed that this scheme would be viable for at least three (3) years, in which time frame, a suitable long term solution would be expected to be deployed.

Note that this plan neither requires nor assumes that already assigned addresses will be reassigned, though if doing so were possible, it would further reduce routing table sizes. It is assumed that routing technology will be capable of dealing with the current routing table size and with some reasonably-small rate of growth. The emphasis of this plan is on significantly slowing the rate of this growth.

This scheme will not affect the deployment of any specific long term plan, and therefore, this document will not discuss any long term

plans for routing and address architectures.

2. Scheme Plan

There are two basic components of this addressing and routing scheme: one, to distribute the allocation of Internet address space and two, to provide a mechanism for the aggregation of routing information.

2.1. Aggregation and its limitations

One major goal of this addressing plan is to allocate Internet address space in such a manner as to allow aggregation of routing information along topological lines. For simple, single-homed clients, the allocation of their address space out of a service provider's space will accomplish this automatically - rather than advertise a separate route for each such client, the service provider may advertise a single, aggregate, route which describes all of the destinations contained within it. Unfortunately, not all sites are singly-connected to the network, so some loss of ability to aggregate is realized for the non simple cases.

There are two situations that cause a loss of aggregation efficiency.

- o Organizations which are multi-homed. Because multi-homed organizations must be advertised into the system by each of their service providers, it is often not feasible to aggregate their routing information into the address space any one of those providers. Note that they still may receive their address allocation out of a service provider's address space (which has other advantages), but their routing information must still be explicitly advertised by most of their service providers (the exception being that if the site's allocation comes out of its least-preferable service provider, then that service provider need not advertise the explicit route - longest-match will insure that its aggregated route is used to get to the site on a non-primary basis). For this reason, the routing cost for these organizations will typically be about the same as it is today.
- o Organizations which move from one service provider to another. This has the effect of "punching a hole" in the aggregation of the original service provider's advertisement. This plan will handle the situation by requiring the newer service provider to advertise a specific advertisement for the new client, which is preferred by virtue of being the longest match. To maintain efficiency of aggregation, it is recommended that organizations which do change service providers plan to eventually migrate their address assignments from the old provider's space to that of the new provider. To this end, it is recommended that mechanisms to facilitate such migration, including improved protocols and procedures for dynamic host address assignment, be developed.

Note that some aggregation efficiency gain can still be had for multi-homed sites (and, in general, for any site composed of multiple, logical IP network numbers) - by allocating a contiguous block of network numbers to the client (as opposed to multiple, independently represented network numbers) the client's routing information may be aggregated into a single (net, mask) pair. Also, since the routing cost associated with assigning a multi-homed site out of a service provider's address space is no greater than the current method of a random allocation by a central authority, it makes sense to allocate all address space out of blocks assigned to

service providers.

It is also worthwhile to mention that since aggregation may occur at multiple levels in the system, it may still be possible to aggregate these anomalous routes at higher levels of whatever hierarchy may be present. For example, if a site is multi-homed to two NSFNet regional networks both of whom obtain their address space from the NSFNet, then aggregation by the NSFNet of routes from the regionals will include all routes to the multi-homed site.

Finally, it should also be noted that deployment of the new addressing plan described in this document may (and should) begin almost immediately but effective use of the plan to aggregate routing information will require changes to some Inter-Domain routing protocols. Likewise, deploying the supernet-capable Inter-Domain protocols without deployment of the new address plan will not allow useful aggregation to occur (in other words, the

addressing plan and routing protocol changes are both required for supernetting, and its resulting reduction in table growth, to be effective.) Note, however, that during the period of time between deployment of the addressing plan and deployment of the new protocols, the size of routing tables may temporarily grow very rapidly. This must be considered when planning the deployment of the two plans.

Note: in the discussion and examples which follow, the network+mask notation is used to represent routing destinations. This is used for illustration only and does not require that routing protocols use this representation in their updates.

2.2. Distributed allocation of address space

The basic idea of the plan is to allocate one or more blocks of Class-C network numbers to each network service provider. Organizations using the network service provider for Internet connectivity are allocated bitmask-oriented subsets of the provider's address space as required.

Note that in contrast to a previously described scheme of subnetting a class-A network number, this plan should not require difficult host changes to work around domain system limitations - since each sub-allocated piece of the address space looks like a class-C network number, delegation of authority for the IN-ADDR.ARPA domain works much the same as it does today - there will just be a lot of class-C network numbers whose IN-ADDR.ARPA delegations all point to the same servers (the same will be true of the root delegating a large block of class-Cs to the network provider, unless the delegation just happens to fall on a byte boundary). It is also the case that this method of aggregating class-C's is somewhat easier to deploy, since it does not require the ability to split a class-A across a routing domain boundary (i.e., non-contiguous subnets).

It is also worthy to mention that once Inter-Domain protocols which support classless network destinations are widely deployed, the rules described by the "supernetting" plan generalize to permit arbitrary super/subnetting of the remaining class-A and class-B address space (the assumption being that classless Inter-Domain protocols will either allow for non-contiguous subnets to exist in the system or that all components of a sub-allocated class-A/B will be contained within a single routing domain). This will allow this plan to continue to be used in the event that the class-C space is exhausted before implementation of a long-term solution is deployed

(there may, however, be further implementation considerations before doing this).

Hierarchical sub-allocation of addresses in this manner implies that clients with addresses allocated out of a given service provider are, for routing purposes, part of that service provider and will be routed via its infrastructure. This implies that routing information about multi-homed organizations, i.e., organizations connected to more than one network service provider, will still need to be known by higher levels in the hierarchy.

The advantages of hierarchical assignment in this fashion are

- a) It is expected to be easier for a relatively small number of service providers to obtain addresses from the central authority, rather than a much larger, and monotonically increasing, number of individual clients. This is not to be considered as a loss of part of the service providers' address space.
- b) Given the current growth of the Internet, a scalable and delegatable method of future allocation of network numbers has to be achieved.

For these reasons, and in the interest of providing a consistent procedure for obtaining Internet addresses, it is recommended that most, if not all, network numbers be distributed through service providers.

3. Cost-benefit analysis

This new method of assigning address through service providers can be put into effect immediately and will, from the start, have the benefit of distributing the currently centralized process of assigning new addresses. Unfortunately, before the benefit of reducing the size of globally-known routing destinations can be achieved, it will be necessary to deploy an Inter-Domain routing protocol capable of handling arbitrary network+mask pairs. Only then will it be possible to aggregate individual class-C networks into larger blocks represented by single routing table entries.

This means that upon introduction, the new addressing plan will not in and of itself help solve the routing table size problem. Once the new Inter-Domain routing protocol is deployed, however, an immediate drop in the number of destinations which clients of the new protocol must carry will occur. A detailed analysis of the magnitude of this expected drop and the permanent reduction in rate of growth is given in the next section.

It should also be noted that the present method of flat address allocations imposes a large bureaucratic cost on the central address

allocation authority. For scaling reasons unrelated to address space exhaustion or routing table overflow, this should be changed. Using the mechanism proposed in this paper will have the happy side effect of distributing the address allocation procedure, greatly reducing the load on the central authority.

3.1. Present Allocation Figures

A back-of-the-envelope analysis of "network-contacts.txt" (available from the DDN NIC) indicates that as of 2/25/92, 46 of 126 class-A network numbers have been allocated (leaving 81) and 5467 of 16256 class-B numbers have been allocated, leaving 10789.

Assuming that recent trends continue, the number of allocated class-B's will continue to double approximately once a year. At this rate of growth, all class-B's will be exhausted within about 15 months.

3.2. Historic growth rates

MM/YY	ROUTES ADVERTISED	MM/YY	ROUTES ADVERTISED
Feb-92	4775	Apr-90	1525
Jan-92	4526	Mar-90	1038
Dec-91	4305	Feb-90	997
Nov-91	3751	Jan-90	927
Oct-91	3556	Dec-89	897
Sep-91	3389	Nov-89	837
Aug-91	3258	Oct-89	809
Jul-91	3086	Sep-89	745
Jun-91	2982	Aug-89	650
May-91	2763	Jul-89	603
Apr-91	2622	Jun-89	564
Mar-91	2501	May-89	516
Feb-91	2417	Apr-89	467
Jan-91	2338	Mar-89	410
Dec-90	2190	Feb-89	384
Nov-90	2125	Jan-89	346
Oct-90	2063	Dec-88	334
Sep-90	1988	Nov-88	313
Aug-90	1894	Oct-88	291
Jul-90	1727	Sep-88	244
Jun-90	1639	Aug-88	217
May-90	1580	Jul-88	173

Table I : Growth in routing table size, total numbers
Source for the routing table size data is MERIT

3.3. Detailed Analysis

There is no technical cost and minimal administrative cost associated with deployment of the new address assignment plan. The administrative cost is basically that of convincing the NIC, the IANA, and the network service providers to agree to this plan, which is not expected to be too difficult. In addition, administrative cost for the central numbering authorities (the NIC and the IANA) will be greatly decreased by the deployment of this plan. To take advantage of aggregation of routing information, however, it is necessary that the capability to represent routes as arbitrary network+mask fields (as opposed to the current class-A/B/C distinction) be added to the common Internet inter-domain routing protocol(s).

3.3.1. Benefits of the new addressing plan

There are two benefits to be had by deploying this plan:

- o The current problem with depletion of the available class-B address space can be ameliorated by assigning more-appropriately sized blocks of class-C's to mid-sized organizations (in the 200-4000 host range).
- o When the improved inter-domain routing protocol is deployed, an immediate decrease in the number routing table entries followed by a significant reduction in the rate growth of routing table size should occur (for default-free routers).

approximately 6%. This is in clear contrast to the current annual growth of 150%. This analysis also assumes an immediate deployment of this plan with full compliance. Note that this analysis assumes only a single level of route aggregation in the current Internet - intelligent address allocation should significantly improve this.

Clearly, this is not a very conservative assumption in the Internet environment nor can 100% adoption of this proposal be expected. Still, with only a 90% participation in this proposal by service providers, at the end of the target three years, global routing table size will be "only" $4700 + 800 + 145 + 7500 = 13145$ routes -- without any action, the routing table will grow to approximately 75000 routes during that time period.

4. Changes to Inter-Domain routing protocols

In order to support supernetting efficiently, it is clear that some changes will need to be made to both routing protocols themselves and to the way in which routing information is interpreted. In the case of "new" inter-domain protocols, the actual protocol syntax changes should be relatively minor. This mechanism will not work with older inter-domain protocols such as EGP2; the only ways to interoperate with old systems using such protocols are either to use existing mechanisms for providing "default" routes or b) require that new routers talking to old routers "explode" supernet information into individual network numbers. Since the first of these is trivial while the latter is cumbersome (at best -- consider the memory requirements it imposes on the receiver of the exploded information), it is recommended that the first approach be used -- that older systems to continue to the mechanisms they currently employ for default handling.

Note that a basic assumption of this plan is that those organizations which need to import "supernet" information into their routing systems must run IGPs (such as OSPF[RFC1267]) which support classless routes. Systems running older IGPs may still advertise and receive "supernet" information, but they will not be able to propagate such information through their routing domains.

4.1. Protocol-independent semantic changes

There are two fundamental changes which must be applied to Inter-Domain routing protocols in order for this plan to work. First, the concept of network "class" needs to be deprecated - this plan assumes that routing destinations are represented by network+mask pairs and that routing is done on a longest-match basis (i.e., for a given destination which matches multiple network+mask pairs, the match with the longest mask is used). Second, current Inter-Domain protocols generally do not support the concept of route aggregation, so the new semantics need to be implemented mechanisms that routers use to interpret routing information returned by the Inter-Domain protocols. In particular, when doing aggregation, dealing with multi-homed sites or destinations which change service providers is difficult. Fortunately, it is possible to define several fairly simple rules for dealing with such cases.

4.2. Rules for route advertisement

1. Routing to all destinations must be done on a longest-match basis only. This implies that destinations which are multi-homed relative to a routing domain must always be explicitly announced into that routing domain - they cannot be summarized

the class-A/B/C classification of network numbers.

4.3. How the rules work

Rule #1 guarantees that the routing algorithm used is consistent across implementations and consistent with other routing protocols, such as OSPF. Multi-homed networks are always explicitly advertised by every service provider through which they are routed even if they are a specific subset of one service provider's aggregate (if they are not, they clearly must be explicitly advertised). It may seem as if the "primary" service provider could advertise the multi-homed site implicitly as part of its aggregate, but the assumption that longest-match routing is always done causes this not to work.

Rule #2 guarantees that no routing loops form due to aggregation. Consider a mid-level network which has been allocated the 2048 class-C networks starting with 192.24.0.0 (see the example in section 5 for more on this). The mid-level advertises to a "backbone" 192.24.0.0/255.248.0.0. Assume that the "backbone", in turn, has been allocated the block of networks 192.0.0.0/255.0.0.0. The backbone will then advertise this aggregate route to the mid-level. Now, if the mid-level loses internal connectivity to the network 192.24.1.0/255.255.255.0 (which is part of its aggregate), traffic from the "backbone" to the mid-level to destination 192.24.1.1 will follow the mid-level's advertised route. When that traffic gets to the mid-level, however, the mid-level *must not* follow the route

192.0.0.0/255.0.0.0 it learned from the backbone, since that would result in a routing loop. Rule #2 says that the mid-level may not follow a less-specific route for a destination which matches one of its own aggregated routes. Note that handling of the "default" route (0.0.0.0/0.0.0.0) is a special case of this rule - a network must not follow the default to destinations which are part of one of its aggregated advertisements.

4.4. Responsibility for and configuration of aggregation

The AS which owns a range of addresses has the sole authority for aggregation of its address space. In the usual case, the AS will install manual configuration commands in its border routers to aggregate some portion of its address space. As AS can also delegate aggregation authority to another AS. In this case, aggregation is done in the other AS by one of its border routers.

When an inter-domain border router performs route aggregation, it needs to know the range of the block of IP addresses to be aggregated. The basic principle is that it should aggregate as much as possible but not to aggregate those routes which cannot be treated as part of a single unit due to multi-homing, policy, or other constraints.

One mechanism is to do aggregation solely based on dynamically learned routing information. This has the danger of not specifying a precise enough range since when a route is not present, it is not always possible to distinguish whether it is temporarily unreachable or that it does not belong in the aggregate. Purely dynamic routing also does not allow the flexibility of defining what to aggregate within a range. The other mechanism is to do all aggregation based on ranges of blocks of IP addresses preconfigured in the router. It is recommended that preconfiguration be used, since it more flexible and allows precise specification of the range of destinations to aggregate.

Preconfiguration does require some manually-maintained configuration

information, but not excessively more so than what router administrators already maintain today. As an addition to the amount of information that must be typed in and maintained by a human, preconfiguration is just a line or two defining the range of the block of IP addresses to aggregate. In terms of gathering the information, if the advertising router is doing the aggregation, its administrator knows the information because the aggregation ranges are assigned to its domain. If the receiving domain has been granted the authority to and task of performing aggregation, the information would be known as part of the agreement to delegate aggregation. Given that it is common practice that a network administrator learns

from its neighbor which routes it should be willing to accept, preconfiguration of aggregation information does not introduce additional administrative overhead.

5. Example of new allocation and routing

5.1. Address allocation

Consider the block of 2048 class-C network numbers beginning with 192.24.0.0 (0xC0180000 and ending with 192.31.255.0 (0xC01FFF00) allocated to a single network provider, "RA". A "supernetted" route to this block of network numbers would be described as 192.24.0.0 with mask of 255.248.0.0 (0xFFF80000).

Assume this service provider connects six clients in the following order (significant because it demonstrates how temporary "holes" may form in the service provider's address space):

- "C1" requiring fewer than 2048 addresses (8 class-C networks)
- "C2" requiring fewer than 4096 addresses (16 class-C networks)
- "C3" requiring fewer than 1024 addresses (4 class-C networks)
- "C4" requiring fewer than 1024 addresses (4 class-C networks)
- "C5" requiring fewer than 512 addresses (2 class-C networks)
- "C6" requiring fewer than 512 addresses (2 class-C networks)

In all cases, the number of IP addresses "required" by each client is assumed to allow for significant growth. The service provider allocates its address space as follows:

- C1: allocate 192.24.0 through 192.24.7. This block of networks is described by the "supernet" route 192.24.0.0 and mask 255.255.248.0
- C2: allocate 192.24.16 through 192.24.31. This block is described by the route 192.24.16.0, mask 255.255.240.0
- C3: allocate 192.24.8 through 192.24.11. This block is described by the route 192.24.8.0, mask 255.255.252.0
- C4: allocate 192.24.12 through 192.24.15. This block is described by the route 192.24.12.0, mask 255.255.252.0
- C5: allocate 192.24.32 and 192.24.33. This block is described by the route 192.24.32.0, mask 255.255.254.0
- C6: allocate 192.24.34 and 192.24.35. This block is described by

the route 192.24.34.0, mask 255.255.254.0

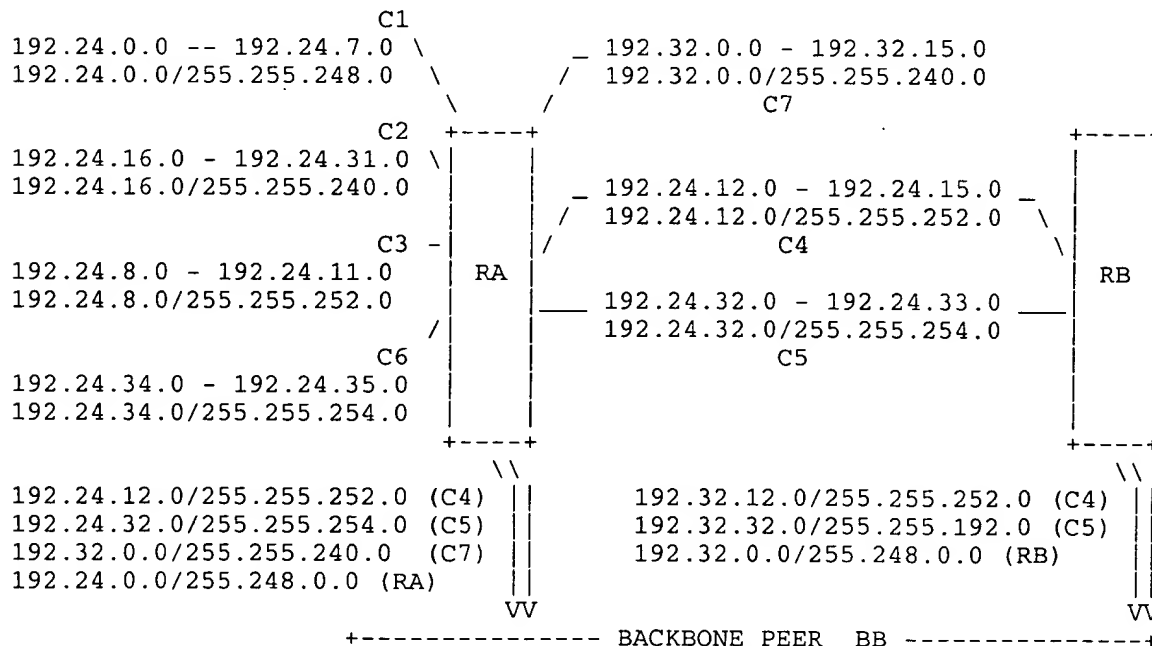
Note that if the network provider uses an IGP which can support classless networks, he can (but doesn't have to) perform "supernetting" at the point where he connects to his clients and therefore only maintain six distinct routes for the 36 class-C network numbers. If not, explicit routes to all 36 class-C network numbers will have to be carried by the IGP.

To make this example more realistic, assume that C4 and C5 are multi-homed through some other service provider, "RB". Further assume the existence of a client "C7" which was originally connected to "RB" but has moved to "RA". For this reason, it has a block of network numbers which are allocated out "RB"'s block of (the next) 2048 class-C network numbers:

C7: allocate 192.32.0 through 192.32.15. This block is described by the route 192.32.0, mask 255.255.240.0

For the multi-homed clients, we will assume that C4 is advertised as primary via "RA" and secondary via "RB"; C5 is primary via "RB" and secondary via "RA". To connect this mess together, we will assume that "RA" and "RB" are connected via some common "backbone" provider "BB".

Graphically, this simple topology looks something like this:



5.2. Routing advertisements

To follow rule #1, RA will need to advertise the block of addresses that it was given and C7. Since C4 and C5 are multi-homed, they must also be advertised.

Advertisements from "RA" to "BB" will be:

192.24.12.0/255.255.252.0 primary	(advertises C4)
192.24.32.0/255.255.254.0 secondary	(advertises C5)
192.32.0.0/255.255.240.0 primary	(advertises C7)
192.24.0.0/255.248.0.0 primary	(advertises remainder of RA)

For RB, the advertisements must also include C4 and C5 as well as it's block of addresses. Further, RB may advertise that C7 is unreachable.

Advertisements from "RB" to "BB" will be:

192.24.12.0/255.255.252.0	secondary	(advertises C4)
192.24.32.0/255.255.254.0	primary	(advertises C5)
192.32.0.0/255.248.0.0	primary	(advertises remainder of RB)

To illustrate the problem alluded to by the "note" in section 4.2, consider what happens if RA loses connectivity to C7 (the client which is allocated out of RB's space). In a stateful protocol, RA will announce to BB that 192.32.0.0/255.255.240.0 has become unreachable. Now, when BB flushes this information out of its routing table, any future traffic sent through it for this destination will be forwarded to RB (where it will be dropped according to Rule #2) by virtue of RB's less specific match 192.32.0.0/255.248.0.0. While this does not cause an operational problem (C7 is unreachable in any case), it does create some extra traffic across "BB" (and may also prove confusing to a network manager debugging the outage with "traceroute"). A mechanism to cache such unreachability information would help here, but is beyond the scope of this document (such a mechanism is also not implementable in the near-term).

6. Transitioning to a long term solution

This solution does not change the Internet routing and addressing architectures. Hence, transitioning to a more long term solution is not affected by the deployment of this plan.

7. Conclusions

We are all aware of the growth in routing complexity, and the rapid increase in allocation of network numbers. Given the rate at which this growth is being observed, we expect to run out in a few short years.

If the inter-domain routing protocol supports carrying network routes with associated masks, all of the major concerns demonstrated in this paper would be eliminated.

One of the influential factors which permits maximal exploitation of the advantages of this plan is the number of people who agree to use it. It is hoped that having the IAB and the Internet society bless this plan would go a long way in the wide deployment, and hence benefit of this plan.

If service providers start charging networks for advertising network numbers, this would be a very great incentive to share the address space, and hence the associated costs of advertising routes to service providers.

8. Recommendations

The NIC should begin to hand out large blocks of class-C addresses to network service providers. Each block must fall on bit boundaries and should be large enough to serve the provider for two years.

Further, the NIC should distribute very large blocks to continental and national network service organizations to allow additional levels of aggregation to take place at the major backbone networks.

Service providers will further allocate power-of-two blocks of

class-C addresses from their address space to their subscribers.

All organizations, including those which are multi-homed, should obtain address space from their provider (or one of their providers, in the case of the multi-homed). These blocks should also fall on bit boundaries to permit easy route aggregation.

To allow effective use of this new addressing plan to reduce propagated routing information, appropriate IETF WGs will specify the modifications needed to Inter-Domain routing protocols. Implementation and deployment of these modifications should occur as quickly as possible.

9. Bibliography

[RFC1247] Moy, J, "The OSPF Specification Version 2", January 1991.

10. Security Considerations

Security issues are not discussed in this memo.

11. Authors' Addresses

Vince Fuller
BARRNet
Pine Hall 115
Stanford, CA, 94305-4122
email: vaf@Stanford.EDU

Tony Li
cisco Systems, Inc.
1525 O'Brien Drive
Menlo Park, CA 94025
email: tli@cisco.com

Jessica (Jie Yun) Yu
Merit Network, Inc.
1071 Beal Ave.
Ann Arbor, MI 48109
email: jyy@merit.edu

Kannan Varadhan
Internet Engineer, OARnet
1224, Kinnear Road,
Columbus, OH 43212
email: kannan@oar.net

[[Index](#) | [Search](#) | [What's New](#) | [Comments](#) | [Help](#)]

Comments/Questions about this archive ? Send mail to rfc-admin@faqs.org